

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR U.S. LETTERS PATENT

Title:

USING LOCAL STORAGE TO HANDLE MULTIPLE OUTSTANDING REQUESTS IN A
SCI SYSTEM

Inventor:

Bryan Hornung
1108 Creekfield Drive
Plano, Texas 75075
Citizenship: US

Bryan Marietta
13117 Boomer Lane
Austin, Texas 78729
Citizenship: US

Robert K. King
2401 Ohio Drive #1708
Plano, Texas 75093
Citizenship: US

USING LOCAL STORAGE TO HANDLE MULTIPLE OUTSTANDING REQUESTS IN A SCI SYSTEM

RELATED APPLICATION

[0001] This application is a divisional application of U.S. Patent Application No. 08/797,674, filed January 31, 1997, entitled "USING LOCAL STORAGE TO HANDLE MULTIPLE OUTSTANDING REQUESTS IN A SCI SYSTEM," which is incorporated herein by reference.

TECHNICAL FIELD

[0002] This invention relates in general to memory accesses in multi-node, multi-processor, cache coherent non-uniform memory access system and relates in particular to managing multiple requests in such a system.

BACKGROUND

[0003] A Scalable Coherent Interface (SCI) Based System coherency flow requires multiple memory accesses. Each access takes many cycles, and therefore, the entire flow takes a great deal of time. The bandwidth of the SCI based system, designed with only one outstanding request, is determined by the latency of each flow. Even though in this type of system, the wires themselves are rated at gigabytes per second, the actual useful bandwidth for each node is limited to closer to 30 to 40 megabytes per second. The reason for this, is that the existing system has enough resources in the SCI controller to handle only one request or response at a time.

[0004] Therefore, there is a need in the art for a method and system that will use more of the available bandwidth of the system by allowing the system to have more than one outstanding request.

SUMMARY

[0005] This need and others are achieved in a system in which one embodiment has local storage for the cache line and tag, and a Contents Addressable Memory (CAM) for the cache line address, is used in the SCI controller to allow numerous outstanding requests or flows to be active at one time. All responses from the SCI ring that generate new SCI requests are handled in the controller without requiring additional memory accesses from the local memory. All conflicts with other SCI cache requests and outstanding flows are also handled by the controller.

[0006] One technical advantage of the present invention is to use a request activation queue to store a request until there are resources available on the SCI ring to handle the request.

[0007] Another technical advantage of the present invention is to use a response activation queue to hold a pointer to a CAM memory location and a table location, so that when the MAC has the required resources to handle the response, the response packet will be formed from the information in the response activation queue.

[0008] A further technical advantage of the present invention is to use a SCI table to store information identifying which memory locations already have outstanding access requests.

[0009] A further technical advantage of the present invention is to use a content addressable memory with match ports to check if a local or ring request is to access a memory location that already has an outstanding request or response.

[0010] The foregoing has outlined rather broadly the features and technical advantages of the present invention in order that the detailed description of the invention that follows may be better understood. Additional features and advantages of the invention will be described hereinafter which form the subject of the claims of the invention. It should be appreciated by those skilled in the art that the conception and the specific embodiment disclosed may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present invention. It should also be realized by those

skilled in the art that such equivalent constructions do not depart from the spirit and scope of the invention as set forth in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, in which:

[0012] FIGURE 1 shows a single node of a multi-node, multi-processor system that uses the inventive TAC arrangement;

[0013] FIGURE 2 shows high level block diagram of the inventive TAC arrangement; and

[0014] FIGURE 3 shows the SCI table field definitions.

DETAILED DESCRIPTION

[0015] FIGURE 1 depicts a single node of a multi-node, multi-processor computer system. The overall system may have a plurality of the nodes shown in FIGURE 1.

[0016] Each node, in the embodiment shown, can support up to sixteen processors 110. These processors 110 are connected to processor agent chips (PACs) 111. The function of each PAC 111 is to transmit requests from its associated processors 110 through cross bar router chips (RAC) 112 to the memory access chips (MAC) 113 and then forward the responses back to the requesting processor. Each PAC 111 has an input/output (I/O) subsystem 117. Each MAC 113 controls access to its associated coherent memory 114. Each MAC 113 is connected to four banks of memory 114 (only two are shown for simplicity). Each bank of memory has four dual in-line memory module boards (DIMM).

[0017] When a processor 110 generates a request to access memory (or other resource), the associated PAC 111 sends the request through the proper RAC 112 to a MAC 113. If the request is destined for memory 114 on the local node, MAC 113 accesses the memory attached to it. If the request is destined for memory on another node, MAC 113 forwards the request to TAC 115. TAC 115 is the interface between the node and an SCI ring 116. TAC 115 is also known as a toroidal access chip or a SCI controller. The SCI rings 116 interconnect the nodes in the multi-node system.

[0018] FIGURE 2 shows a high level block diagram of the inventive TAC 200. The following describes how data packets flow through this device. In general, the request will come in from a MAC 113 through an interface, MAC-to-TAC Control 201. A request will be split off by MAC-to-TAC control 201 and put into MAC Request In Queue 202. The Table Initialization State Machine 203 receives the requests from queue 202.

[0019] Table initialization state machine 203 will determine the first state in a flow, and then write that information into SCI Table 204. State machine 203 will also write any data that came in with the request into SCI table 204 and then write the address into Address CAM 205. Table initialization machine 203 will then send a request to Request Activation Queue 206. The request will remain in request activation queue 206 until there are sources available on ring 116 to handle the delivery of this request. SCI Request Packet Assembly

207 will translate the request into the symbols necessary to generate an SCI request. That request will be sent out on Datapump 208 to rings 116.

[0020] On a remote node, that request will be sent to a memory or cache 114, and that will generate a response. The response will come in on datapump 208. The response will travel on the SCI Response In wires and be delivered to SCI Response Engine 209. Response engine 209 will then read the contents of the table and the CAM that were written previously and will determine what to do next. The system is able to send another request to ring 116. The system may also send a response to MAC 113. Both the response sent to MAC 113, and the request sent to ring 116 can be busy, so the system has the capability to wait for resources while receiving response from ring 116.

[0021] Therefore, a request to ring 116 will use request activation queue 206, and the response to MAC 113 will use Response Activation Queue 210. As SCI response engine 209 will take the response, read the contents of address CAM 205 for the address, and SCI table 204 for the state, and then use Next Cache State Table 211 to determine what is to be done next.

[0022] If a response is generated and the flow is done, and there are enough response resources to actually generate the response to MAC 113, then engine 209 sends the response packet through MAC Response Out Queue 212 and then through TAC-to-MAC Control 213, which arbitrates with finality between MAC response out queue 212 and MAC Request Out 216 queue, and sends the proper packet to MAC 113.

[0023] As mentioned above, a request will then go out to ring 116. On another, remote node, that request will come to the node from ring 116 through that node's datapump 208. The request will enter the datapump 208, and then be sent to the remote node's SCI Request Engine 214. It will then check the address of that request with all addresses that are currently being worked on in that TAC 115. This check is done by the Contents Addressable Memory or Address CAM 205.

[0024] If there is a hit, the entry number generated by CAM 205 is then used to access SCI table 204 and the request is handled locally, and the response is sent out back to the ring for muxing between local responses from SCI request engine 214 and the MAC

responses by SCI Response Out mux 215. If there was no hit in CAM 205, the request is sent to local MAC 113 to be handled by the memory controller, thus, the request goes into MAC Request Out Queue 216 through TAC-to-MAC Control 213.

[0025] The memory controller then handles that request and sends response back to TAC 115. The response comes in on MAC-to-TAC control 213. That response will then be routed to MAC Response In Queue 217, which will then be checked by Response SCI question block 218, which determines whether the response was generated for one of the local node requests, or if the response was generated from a ring request from a remote node. Since this is response from a ring request, then it is reformatted into ring packets and sent to mux 215 where it will then be forwarded to datapump 218.

[0026] The significant features of this system 200 that allow it to handle many outstanding requests and responses at the same time are Address CAM 205, SCI Table 204, Request Activation Queue 206, and Response Activation Queue 210. In this particular design, both CAM 205 and table 204 can handle 32 different requests at the same time. CAM 205 has within it 32 addresses, and the table 204 contains 32 states and 32 sets of data for any of the lines. Request Activation Queue 206 contains essentially just the pointer to SCI Table 204 and to an address location in CAM 205.

[0027] The SCI Request Packet Assembly 207 uses that pointer from Request Activation Queue 206 to read table 204 and the CAM 205 to assemble a request packet. These request packets can be up to 12 symbols long and are stored in the datapump until they are actually put on the ring.

[0028] For a response, MAC Response Out Queue 212 also holds fully assembled packets. Response Activation Queue 210 also holds a pointer to a CAM 205 location and to a table 204 location. When MAC Response Out Queue 212 has room, SCI Response Engine 209 will take the top response from of Response Activation Queue 210, use that index to read SCI table 204 and the address CAM 205 and will then assemble the response packet at that time.

[0029] As previously stated, CAM 205 is a contents addressable memory. This means is that there are match ports, wherein the data at the match ports can be applied to

simultaneously check every location in CAM 205 to see if data exists that is identical to the data at the match port. If the data is identical, then CAM 205 generates an index which can be used by the various other state machines to access SCI table 204.

[0030] For example, State Machine 203, the table initialization state machine, checks all requests coming in from MAC 113 to see if there is already a request for that address in TAC 115. TAC 115 can only handle only one request for a given address at a time, so table initialization machine 203 will take the address generated by MAC 113 and apply it to CAM 205 with the data and match port, and CAM 205 will return with a hit or miss.

[0031] If there is a hit, CAM 205 will return with an index that table initialization machine 203 can use to access SCI table 204. SCI Response Engine 209 uses the index supplied by the response packet to address CAM 205. SCI Request Engine 214, takes an address that it gets from ring 116 and applies it to CAM 205 and using its match port CAM 205 will return with either a hit or miss. If it is a hit, it will return the index, which SCI Request Engine 214 can then use to access to SCI table 204. Other things that can access CAM 204 are Request Packet Assembly 207 which uses an index stored in Request Activation Queue 206 to read an address.

[0032] SCI Response Engine 209 only uses the read feature of CAM 204. This engine 209 received a transaction ID from the response off ring 116. This transaction ID is the exact same ID that was used to access CAM 205 and table 204 while generating the request by request activation queue 206.

[0033] When table initialization state machine 203 checks CAM 204 for a match on the address it received in a new request from MAC 113, machine 203 will do one of two things, depending on whether there is a hit or a miss. A hit means there is already an outstanding request in TAC 115 for a given address. In this case, the new request is chained onto the back of the other request so that it can be handled sequentially. If there is a miss, which should be the normal case, a new request is immediately generated and sent out to ring 116.

[0034] SCI Request Engine 214, also checks for a hit or miss on CAM 205. In the case when there is a hit, SCI Request Engine 214 handles the request locally with information contained in CAM 205 and table 204, and if there is a miss, the request is forwarded on to MAC 113 for handling by the memory controller.

[0035] SCI Table 204 is a 32 bit entry table that contains information described in FIGURE 3. The table 300, includes a table_state. This state can be unused, which means that the table of this particular entry has not been used. The state can also be queued, which means that this entry is queued behind an active entry. Waiting means that this entry is waiting for more information from a MAC 113 before it can generate a request. Queued Waiting means the entry is queued behind another active request, and when that request is done, it will then have to wait for still more information from a MAC 113 before continuing. Active means that it is in the middle of an active flow, and Done means that the flow is done, but its resources have not been de-allocated.

[0036] Flow_Type 302 contains the transaction type. These are the different transactions that TAC 113 may perform. TAC 113 can perform read shared, read private, read rollout, read current, write purge, global flush, increment update, or various non-coherent transactions.

[0037] Master_ID 303 is the transaction master that was received from MAC 113 and indicates that this was the owner of the original request.

[0038] Transaction_ID 304 is also received from MAC 113, and indicates that this is the particular transaction from a given master. The transaction ID and the Master ID combined together are unique identifiers which allows responses to be returned to the requester.

[0039] The c_state 305 or cache state field is a transient cache state.

[0040] The c_forw 306 or cache forward is an SCI cache forward pointer.

[0041] The c_back 307 or cache backward field is the SCI backward pointer.

[0042] The shared_phase 308 is the shared phase used in the increment update flow.

[0043] The T field 309 encodes the type of access being performed with non-coherent accesses. Non-coherent accesses can go to memory space or they can go to CSR space.

[0044] The next field 310 is the next chained entry. This is used for chaining entries together when there are multiple requests to the same address outstanding in TAC 113.

[0045] The weak bit 311 is used in read private flow to determine whether there are weak or strong ordered responses.

[0046] The magic bit 312 is called magic because it has a number of different functions, depending on the type of flow being done. One major function is that it marks a rollout as a flush. A flush and rollout are identical except a flush sends a response at the end. Another major function is that it specifies that data has been returned for weak ordered flows.

[0047] The rollout phase bits 313 are used to specify additional transient cache states to resolve rollout and increment update collisions.

[0048] Although the present invention and its advantages have been described in detail, it should be understood that various changes, substitutions and alterations can be made herein without departing from the spirit and scope of the invention as defined by the appended claims.